

Collaboration and security in CNL's virtual laboratory

Andrew Tokmakoff¹, Yuri Demchenko² and Martin Snijders¹

¹*Telematica Instituut, P.O. Box 589, 7500 AN Enschede, The Netherlands*

²*Universiteit van Amsterdam, Advanced Internet Research Group, Kruislaan 403,
NL-1098 SJ Amsterdam, The Netherlands*

Andrew.Tokmakoff@telin.nl

demch@science.uva.nl

Martin.Snijders@telin.nl

Abstract

Collaborative and distributed workspaces provide opportunities for organizations to make more efficient use of expensive and specialised laboratory equipment. Along with such opportunities lies a clear need to ensure secure access to, and operation of applications and services. In this paper, we provide an introduction to the CNL Virtual Laboratory, followed by a discussion of its security architecture, with particular emphasis being placed on Authorization of user actions in the collaborative environment.

1 Introduction

The process industry makes extensive use of advanced laboratory equipment, such as electron microscopes, equipment for surface analysis and mass spectrometers. Due to high initial outlay and operational costs, and the expertise required to operate such equipment, laboratories tend not to have all the necessary equipment in-house. The Collaboratory.nl project (CNL) investigates how technologies for remote operation of laboratory equipment can be integrated with existing groupware for enhanced remote collaboration. Such a virtual laboratory offers the same possibilities as a traditional laboratory, but also enables laboratory staff to utilise the equipment and expertise of third parties. CNL's research addresses the following points:

- Industrial research provides companies with a competitive edge. It is therefore extremely important that research infrastructure and data are properly secured.
- The virtual laboratory should not exist in isolation from other scientific computing initiatives. It should take emerging standards into account, such as those in the areas of Grid-based computing and collaborative systems.
- To allow commercial exploitation of such a virtual laboratory, business models and charging mechanisms need to be established to allow cost distribution over the appropriate parties.

In this paper, we elaborate on efforts to address the security of the CNL virtual laboratory and introduce the project's first software deliverable, that being a concept demonstrator of the CNL system.

2 The CNL Virtual Laboratory

In order to demonstrate a commercially operable Virtual Laboratory, CNL has identified the following major functional blocks:

- Remote Access and Use of Instruments (instrument client apps and services),
- Collaboration Toolset (chat, whiteboard, persistent workspace etc),
- Security (user authentication/authorization, information confidentiality and integrity etc),
- Business Management (sample tracking/tracing, job management etc.),
- Financial Exploitation (metering, accounting, charging).

The Virtual Laboratory makes extensive use of Jobs (a key concept), which can be created by Analysts (a privileged set of users). They contain:

- Job processing information (workflow),
- details regarding the sample(s) to be analysed,
- the customer who commissioned the work,
- the instruments (denoted as resources) which will be used within the Job,
- the users (and their role[1]) in the Job.

A set of use-cases for the Virtual Laboratory were developed, which served to set scope and also provided input for the initial Technical Requirements activities. The use-cases identified a number of fundamental usages that allow remote users to interact with an operator (where necessary), to remotely control an instrument, to persistently share information amongst themselves.

Some key requirements for the Virtual Laboratory concerned more pragmatic deployment issues, related to its ability to operate in strictly-controlled corporate computing environments. Such environments typically restrict the applications that may be deployed on desktop machines, mandate the use of proxies for outgoing port 80 traffic, and employ strict (ingoing and outgoing) firewall policies. Some issues related to satisfying these requirements have been addressed, whilst others remain ongoing.

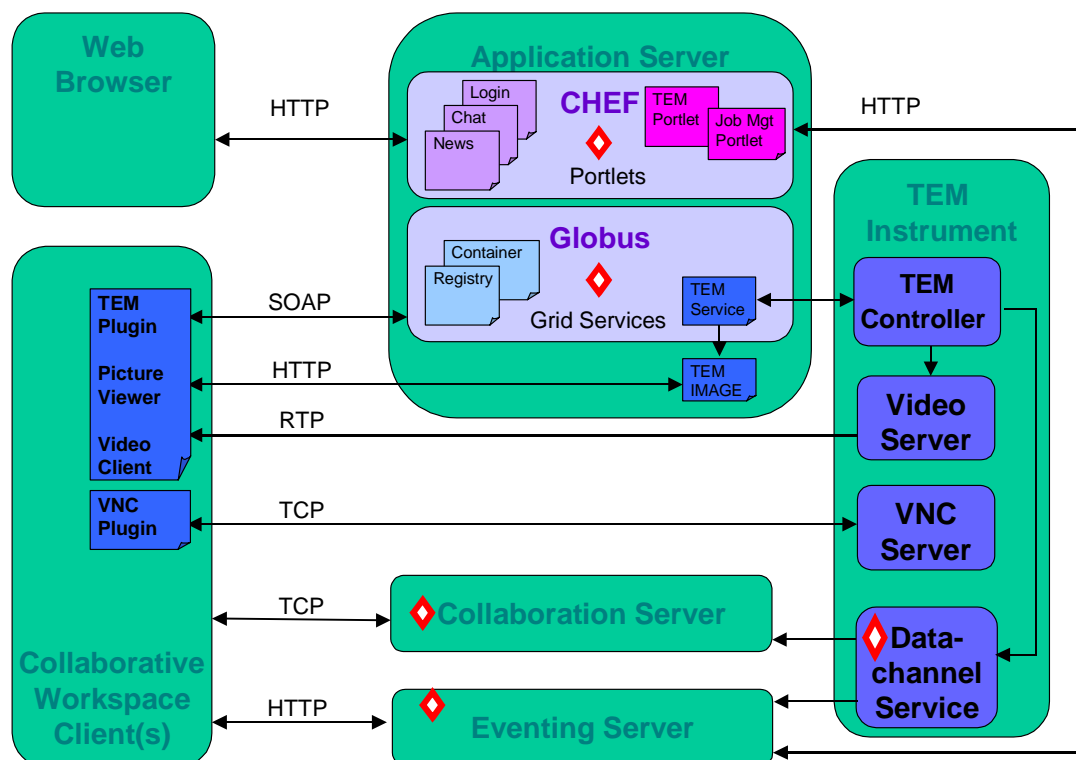


Figure 1: Collaboration and Instrument Control in CNL

2.1 Architecture

The CNL Virtual Laboratory's Architecture (which has been simplified to include only a single class of Instrument) is shown in Figure 1. There are five main sets of services (each denoted by a triangle):

- a set of portals (Portlets) that act as an "entry-point" to the system, a Grid server, which hosts Grid services for-CNLI Instruments,
- a Collaboration server, which allows users of the "Collaborative Workspace Clients" to interact,
- an Eventing Server which provides transient and persistent messaging services
- a Datachannel service, which monitors local disk storage and allows instrument data to enter the collaborative environment's workspace.

Instruments are made available as Grid Services, which are then "connected" to the dedicated instrument control infrastructure at the Instrument site. This approach provides two benefits:

- the instrument control interface can be well-defined using Grid technologies
- access and control to the instrument can be regulated through the use of CNL's security infrastructure (detailed further in Section 4).

It is expected that Instruments will be located in numerous different locations, and it follows that an Instrument's Grid service and its controller would be

deployed at the same physical location, although this is not strictly necessary (there are significant security implications).

On the client side, users log in to the system through the use of a web-browser, and make use of a "rich" collaborative toolset client, which hosts a set of CNL tools (covered in more detail in section 2.2).

The CNL workspace is divided into two main areas: a web-based portal¹ and a client-side workspace², which were both adapted from upon existing open source systems. Together, they provide CNL users with a Portal for administrative tasks (such as sample tracking and tracing, billing etc), and a dynamic collaborative workspace that can host "tools" for remote access and control of CNL resources.

The web-based Portal is a "launch-pad" for users who only need a browser and Java Web Start³ on their desktop machine. Users can then launch the collaborative workspace from within the Portal and commence collaborative work.

2.2 Collaborative Tooling

The CNL collaborative workspace provides a basic set of collaborative functions (chat, shared whiteboard, persistent workspace) which can be used to enable co-

¹ <http://www.chefproject.org>

² <http://surabaya.sourceforge.net>

³ <http://java.sun.com/products/javawebstart/>

operative distributed working, in conjunction with “plain old telephone service” (POTS) voice connectivity.

The CNL system also has a number of Instrument-specific tools that have been built to demonstrate the system’s use with X-Ray Photoelectron Spectroscopy (XPS) and Transmission Electron Microscope (TEM) laboratory instruments. It currently accommodates two instrument interaction tools, each of which takes a different approach towards enabling access and remote-control of CNL resources. This provides some interesting research perspectives with regard to end-user experiences when using laboratory instrument via CNL (future work). These tools can be divided into two main classes:

- **Instrument-specific tools:** provide a dedicated interface towards the instrument/resource. All controls are engineered for the specific type of resource that is to be interacted with. The communication protocols used by the tool for interacting with the instrument is left up to the tool implementation and may be tool-specific (e.g. proprietary control protocols or control via well-specified WSDL interfaces). An example in CNL is the tool that is used to access and control a TEM.
- **Instrument-agnostic tools:** these tools make use of desktop-sharing functionality and can communicate with any instrument/resource that provides an appropriate desktop-sharing server. The use of remote desktop sharing is a very limited form of collaboration and requires additional support for remote file management and access control. An example of this is the tool which is used to access and control a XPS.

Within CNL’s collaborative environment, the availability of tools is always Job-dependent, which is regulated by a Job’s context, which specifies the tool(s) associated with the Job. If a Job requires the use of a certain type of instrument (or resource), then one or more tools that are associated with the resource are automatically activated in the end-user’s collaborative toolset. Basic collaborative tools are always available to the user, irrespective of the job they are currently in. However, such tools also have the ability to take advantage of the job’s context.

For example, when a file is posted to a job’s shared workspace, it only becomes visible to users that are a member of that Job and who are currently working within the Job’s context, even though the shared workspace tool itself is still available to all users. Similarly, the use of the Whiteboard and Chat tools are Job centric, which means that only users that are working within the same Job context can directly collaborate.

Figure 1 includes a TEM Instrument, which makes use of a video server to stream the TEM’s video image (via RTP) to its associated control and visualisation Tool, hosted within the collaborative environment.

The Instrument also hosts a Datachannel service, which is used to “funnel” output data (such as still images and data sets) from the Instrument into the collaborative workspace. This generic tool is equally applicable for use with the XPS and other Instruments.

The openness of the Virtual Laboratory is illustrated by the “open” tool architecture which it has adopted. Using this approach, third parties are able to develop their own tools and add them to the CNL system as a plug-in.

Each CNL collaborative tool implements the CNL tool interface, which prescribes methods for starting and stopping tools and allows specification of a tool’s authorization level. Furthermore, the interface offers “common” CNL services to the tools, such as communication channels and an ability to publish images to the shared whiteboard. In this way, the CNL tool framework is able to offer collaborative services to a wide variety of instrument tools.

3 CNL as a Virtual Organisation

Originating from the OGSA, the concept of a Virtual Organisation (VO) allows task-oriented virtualisation of resources and services. A VO consists of a set of individuals and associated distributed resources [2], and is created via business agreements between participating organisations and individuals, each of which contributes their specific resources (computers, services, people, etc.). The agreement defines all resources and services available to VO members and also conditions upon which these resources and services are provided and can be used. As with many “real” organisations, VO’s may have access to all the basic services required to run a typical organisation, but these services may be “physically” and administratively run by member organisations on behalf of the VO.

The VO concept integrates naturally with the Virtual Laboratory’s Job-centric approach, as VO’s can be defined inside a Job description and hence, can be related to specific experiments. In distributed collaborative environments, VO membership management functionality extends beyond typical enterprise identity management concepts and requires multi-institutional federation of people, resources and services management.

When created as a Grid Service using an OGSI VO Factory, a VO instance can supply a context that can be used to dynamically associate users, resources, policies and agreements when making and processing service requests related to a particular VO.

The Virtual Laboratory plans to implement the OGSA VO concept to manage its distributed and dynamic cooperative environment. VO requires specific management and security services, which have been defined in the OGSA framework [3] and which are also being further developed in some leading Grid projects such as LCG [4] and EGEE [5].

4 Security in the Virtual Laboratory

Collaborative applications require a sophisticated multi-dimensional security infrastructure that should be able to manage secure operation of user applications between multiple administrative/organisational and trust domains. Such a security infrastructure should address both policy driven user access control and task/job based trust management.

Early experiences with the use of PKI technology for authentication, authorization and secure communication in collaborative Grid-based projects were documented in [6]. Development of access control solutions based on a proprietary Community Authorisation service (CAS) [1] has been described in [7], and also using a XACML policy-based access control model in [8].

Security in collaborative environments has become vital with the ongoing emergence of computer Grids and related technologies for resource and user group virtualization. Such activities require a high granularity and a customer-driven approach to policy and role-based access control.

The Virtual Laboratory's Job-centric approach uses the Job description as a semantic document, which is created on the basis of a signed order (business agreement) and contains all the information required to run the analysis, including the Job ID and other attributes, assigned users and roles, and a trust/security anchor(s) in a form of customer and/or CNL digital signature. In general, such a Job-centric approach allows binding security services and policies to a particular job or resource.

4.1 Security Architecture

The Virtual Laboratory makes use of a Security Architecture that builds upon emerging WS-Security [9] and OGSA [1] Security services and standards, along with a generic authorisation framework [10]. In the Web Services Architecture, a Web Service can have security services and components added to the service description by applying the WS-Security set of standards. WS-Security components are included within the Globus Toolkit v3.2, and CNL leverages these technologies to achieve a flexible customer-controlled security infrastructure/ environment, which is also generally applicable for collaborative environments. The Virtual Laboratory's Security Architecture includes:

- 1) Communication/transport Security Layer: defines network infrastructure security and uses network security services such as SSL/TLS, IPsec and VPN.
- 2) Messaging Security Layer: based on currently well-defined and supported Web Services platforms SOAP/WS-Security [11], and also makes use of XML Security mechanisms such as XML Signature, XML Encryption, and SAML as a security token exchange format [12].

- 3) Policy Expression and Exchange Layer: defines a set of policies, which can be applied to CNL users when they interact with the environment, which are necessary to ensure multi-domain and multiplatform compatibility.
- 4) Services/Operational Layer defines security services/mechanisms for secure operation of the CNL System in an open environment and includes:
 - Authentication and Identity Management
 - Authorization and Access Control
 - Secure Context Management
 - Auditing/Logging and Notarisation

CNL provides basic security services: authentication and identity management, authorization, information and data confidentiality and integrity, non-repudiation and privacy. It also provides a Single-Sign-On (SSO) facility.

4.2 Authorization of Actions

The authorization (or access control) system uses a policy and role based access control approach, which allows flexible dynamic user access management (see Figure 2). The CNL Authorization (AuthZ) infrastructure consists of:

- a Policy Enforcement Point (PEP) which provides resource-specific authorization decisions, request/response handling and policy-defined obligations enforcement,
- a Policy Decision Point (PDP)/Rule-Based Engine (RBE) [10] that is the central policy-based decision making point. It evaluates authorization requests against the policy defined for a particular job, resource and user attributes/roles.
- a Policy Authority Point (PAP) that is a (potentially distributed) policy store,

To gain access to a specific resource, the Resource Agent makes a request on the PEP for an authorization decision from the Policy Decision Point (PDP), which evaluates the authorization request against the policy defined for a particular job, resource and user attributes/roles. The access policy is defined by the resource owner and/or CNL administrator and stored in the PAP's policy database.

As an example, when a tool is instructed by an End-user to initiate an action (such as a remote-control action on an instrument), the Virtual Laboratory's security sub-system is accessed and requested for authorization of the action. Similarly, whether a tool is even included in an End-user's collaborative toolset interface or not depends upon the End-user's role. For example, certain controls that are restricted for use only by a certain class of user (such as an Analyst) will be disabled for other End-users.

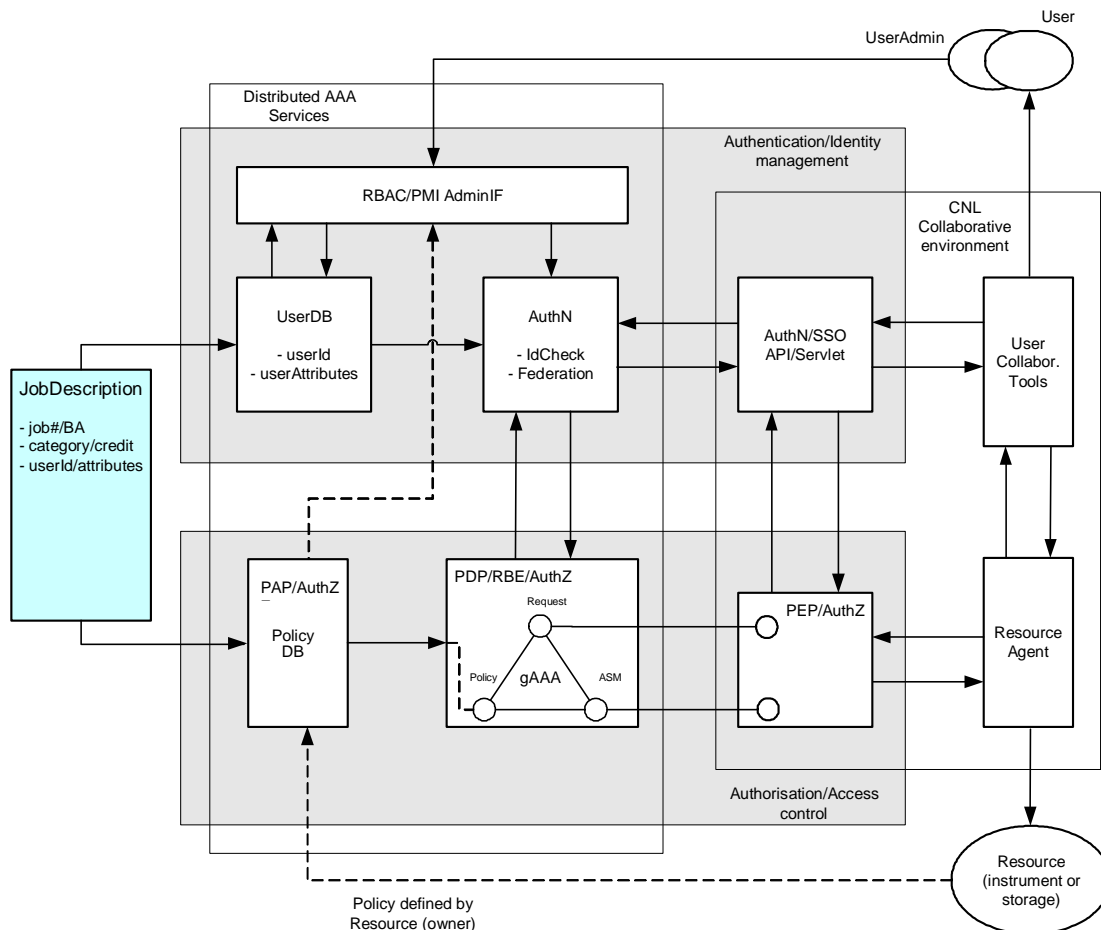


Figure 2: Basic Authentication/Authorization functionality in CNL.

5 Summary

The security approach presented in this paper illustrates the use of Web Services security mechanisms/technologies, to combine the flexibility of the generic AAA Architecture with the XACML policy-based access control model, allowing fine-grained access control and cross-organisation identity management using the VO concept.

CNL represents a typical use case for general Web Services and the OGSA Security framework/architecture. CNL is being developed in close liaison with ongoing Grid related projects including VL-E [13], LCG and EGEE. Maintaining close ties with these projects ensures CNL's future compatibility with emerging Grid Infrastructure located within Europe and throughout the world.

6 Conclusions and Further Work

CNL has now completed its first project year. Within this time frame, we have defined a set of use-cases that sketch a set of activities that the virtual laboratory should be able to support, and created a concept demonstrator that contains many of the virtual laboratory's intended functionalities.

In co-operation with the project's industrial partners, we are now entering a validation phase, and continue

to implement/integrate the Virtual Laboratory's security infrastructure. In addition, we are conducting further researching into business aspects of the virtual laboratory, including topics such as Service Level Agreements (SLA's), job submission and management, sample tracking, and metering/billing of service usage [14].

Acknowledgements

The authors would like to thank Rene van Buuren for timely and constructive comments on drafts of this paper. The results presented in this paper originate from work conducted within the Collaboratory.nl project, a research initiative that explores the possibilities of remote control and use of advanced lab facilities in a distributed and collaborative industrial research setting. The Collaboratory.nl consortium consists of DSM, Philips, Corus, FEI, Telematica Instituut and the University of Amsterdam.

References

- [1] Role Based Access Control (RBAC) – NIST, April 2003. - <http://csrc.nist.gov/rbac>

- [2] Yu. Demchenko, "Virtual Organisations in Computer Grids and Identity Management", Information Security Technical Report - Volume 9, Issue 1, January-March 2004, Pages 59-76.
- [3] The Open Grid Services Architecture, Version 1.0, July 12, 2004,
<http://www.gridforum.org/Meetings/GGF12/Documents/draft-ggf-ogsa-specv1.pdf>
- [4] The LHC Computing Grid Project (LCG) -
<http://lcg.web.cern.ch/LCG/>
- [5] The Enabling Grids for E-science in Europe, or EGEE project - <http://public.eu-egee.org/>
- [6] D. Agarwal, K. Jackson, M. Thompson, "Securing Collaborative Environments", Proceedings of the Workshop on Advanced Collaborative Environments, Edinburgh, Scotland, July 2002.
- [7] D. Agarwal, M. Lorch, M. Thompson, M. Perry, "A New Security Model for Collaborative Environments", Proceedings of the Workshop on Advanced Collaborative Environments, Seattle, WA, June 2003.
- [8] Markus Lorch, Dennis Kafura, Sumit Shah, "An XACML-based Policy Management and Authorization Service for Globus Resources", Work in progress paper, Proc. 4th Int. Workshop on Grid Computing - Grid 2003, 17 November 2003 in Phoenix, AR, USA
- [9] "Security in a Web Services World: A Proposed Architecture and Roadmap", Version 1.0, A joint security whitepaper from IBM Corporation and Microsoft Corporation. April 7, 2002,
<http://www-106.ibm.com/developerworks/library/ws-secmap>
- [10] RFC 2903, Experimental, "Generic AAA Architecture", C. de Laat, G. Gross, L. Gommans, J. Vollbrecht, D. Spence, August 2000,
<ftp://ftp.isi.edu/in-notes/rfc2903.txt>
- [11] OASIS Web Services Security Framework
- [12] Security Assertion Markup Language (SAML) v1.0 - OASIS Standard, November 5, 2002.
- [13] Virtual Laboratory for e-Science - <http://www.vl-e.nl/>
- [14] H. Jonkers (ed.), S.C. Hille, A. Tokmakoff & M. Wibbels, "A functional architecture for the financial exploitation of network-based services", 2001,
<https://doc.telin.nl/dscgi/ds.py/Get/File-13196>